UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/808,232	03/24/2004	Ofir Zohar	ASSIA 21.062	5004

26304 7590 04/20/2007

KATTEN MUCHIN ROSENMAN LLP
575 MADISON AVENUE
NEW YORK, NY 10022-2585

EXAMINER	
FARROKH, HASHEM	
ART UNIT	PAPER NUMBER
2187	

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
3 MONTHS	04/20/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary	Application No.	Applicant(s)	
	10/808,232	ZOHAR ET AL.	
	Examiner	Art Unit	
	Hashem Farrokh	2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *29 January 2007*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-42* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-3,5-22 and 24-42* is/are rejected.

7)☒ Claim(s) *4 and 23* is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☒ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *24 March 2004* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

This Office Action is in response to the Applicant's Remark filed on 1/29/07. There are 42 claims pending in the application; no claims have been amended or canceled.

INFORMATION CONCERNING CLAIMS:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

Claims 1-3, 8-9, 11-13,15, 18, 20-22, 27-28, 30-32, 34-35, 37, and 39-42 are rejected under 35 U.S.C. 102(b) as being anticipated by U.S. Patent No. 6,405,284 B1 to Bridge.

1. In regard to claim 1 Bridge teaches:

A method for storing data **(e.g., abstract)**, comprising:

"distributing a first plurality of groups of logical addresses among one or more storage devices;" **(e.g., column 30, lines 50-54).**

"receiving a second plurality of data-sets containing the data to be stored;" **(e.g., column 27, line 35).**

"assigning each data-set among the plurality of data-sets a number chosen from a first plurality of different numbers;" **(e.g., column 6, lines 47-51; Data Extents 124-128 in Fig. 1).**

"partitioning each data-set into multiple partitions, so that each partition among the multiple partitions receives a sequential partition number;" **(e.g., column 6, lines 42-56).**

"assigning each partition within each data-set to be stored at a specific group of logical addresses in accordance with the sequential partition number of the partition and the number assigned to the data-set;" **(e.g., column 8, lines 62-67; column 10, lines 44-52; Fig. 4).**

"and storing each partition at the assigned specific group of logical addresses." **(e.g., column 4, lines 60-63).**

2. *In regard to claim 8 Bridge teaches:*

"A method for data distribution **(e.g., abstract)**, comprising: receiving at least part of a data-set containing data;" **(e.g., column 27, line 35).**

"delineating the data into multiple partitions;" **(e.g., column 4, lines 50-54).**

"distributing logical addresses among an initial set of storage devices so as to provide a balanced access to the devices;" **(e.g., column 30, lines 50-54).**

"transferring the partitions to the storage devices in accordance with the logical addresses;" **(e.g., column 7, lines 50-55).**

"adding an additional storage device to the initial set **(e.g., column 22, lines 29-32; Fig. 16)**, thus forming an extended set of the storage devices comprising the initial set and the additional storage device;" **(e.g., column 31, lines 1-5; claim 32).**

"and redistributing the logical addresses among the storage devices in the extended set so as to cause a portion of the logical addresses and the partitions stored thereat to be transferred from the storage devices in the initial set to the additional storage device **(e.g., column 31, lines 3-5)**, while maintaining the balanced access and without requiring a substantial transfer of the logical addresses among the storage devices in the initial set." **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

3. In regard to claim 12 Bridge teaches:

"A method for data distribution **(e.g., abstract)**, comprising:"

"receiving at least part of a data-set containing data;" **(e.g., column 27, line 35).**

"delineating the data into multiple partitions;" **(e.g., column 4, lines 50-54).**

"distributing logical addresses among an initial set of storage devices so as to provide a balanced access to the devices;" **(e.g., column 30, lines 50-54).**

"transferring the partitions to the storage devices in accordance with the logical addresses;" **(e.g., column 7, lines 50-55).**

"removing a surplus storage device from the initial set, thus forming a depleted set of the storage devices comprising the initial set less the surplus storage device;" **(e.g., column 5, lines 6-7).**

"and redistributing the logical addresses among the storage devices in the depleted set so as to cause the logical addresses of the surplus device and the partitions stored thereat to be transferred to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of the logical addresses among the storage devices in the depleted set." **(e.g., column 30, lines 60-67 to column 31, lines 1-5).** *Bridge teaches that when a disk drive removed (or added), it changes the configuration of the storage system. Data is dynamically moved to remaining or depleted storage devices to provide load balancing (e.g., see column 3, lines 16-18).*

4. In regard to claim 16 Bridge teaches:

"A data storage system **(e.g., Fig. 1)**, comprising: one or more mass-storage devices **(e.g., Disk Group 104 in Fig. 1)**, coupled to store partitions of data at respective first ranges of logical addresses (LAs);" **(e.g., column 5, 56-62; lines column 9, lines 36-43; Fig. 3).** For example directories in each "root disk group" contain information including logical volume addresses and pointer that identifies the extents or partitions.

"a plurality of interim devices **(e.g. Disk Drives 106-114 in Fig. 1)**, configured to operate independently of one another **(e.g., column 12, lines 12-14**, each interim device being assigned a respective second range of the LAs and coupled to receive the partitions of data from and provide the partitions of data to the one or more mass-

storage devices having LAs within the respective second range;" **(e.g., column 9, lines 4-10; extent B in disk drive 214 in Figs. 2A-2B).** The logical volume or the logical address space is portioned to a plurality of extends. For example extents A-D addresses the stripes in disk drive 214-218 in Fig. 2A.

"and one or more interfaces **(e.g., column 6, lines 20-22; Fig. 1)**, which are adapted to receive input/output (IO) requests from host processors **(e.g., column 5, lines 33-34)**, to identify specified partitions of data in response to the IO requests, to convert the IO requests to converted-IO-requests directed to specified LAs in response to the specified partitions of data, and to direct all the converted-IO-requests to the interim device to which the specified LAs are assigned." **(e.g., column 5, lines 56-62; column 9, lines 36-38).**

5. In regard to claim 20 Bridge teaches:

"A data storage system **(e.g., Fig. 19)**, comprising: one or more storage devices wherein are distributed a first plurality of groups of logical addresses;" **(e.g., column 30, lines 50-54).**

"and a processing unit which is adapted to: receive a second plurality of data-sets containing the data to be stored **(e.g., column 27, line 35)**, assign each data-set among the plurality of data-sets a number chosen from a first plurality of different numbers, partition each data-set into multiple partitions **(e.g., column 4, lines 50-54)**, so that each partition among the multiple partitions receives a sequential partition number **(e.g., Extents A-D and Stripes A0:A3-D0:D3 in Figs 2A-2B)**, assign each partition within

each data-set to be stored at a specific group of logical addresses in the one or more storage devices in accordance with the sequential partition number of the partition and the number assigned to the data-set, and store each partition in the one or more storage devices at the assigned specific group of logical addresses." **(e.g., column 9, lines 4-17).**

6. *In regard to claim 27 Bridge teaches:*

"Data distribution apparatus **(e.g., Fig. 19)**, comprising: an initial set of storage devices among which are distributed logical addresses so as to provide a balanced access to the devices;" **(e.g., column 23, lines 49-55).**

"an additional storage device to the initial set **(e.g., column 22, lines 29-32; Fig. 16)**, thus forming an extended set of the storage devices comprising the initial set and the additional storage device;" **(e.g., column 31, lines 1-5; claim 32).**

"and a processor **(e.g., Processor 1904 in Fig. 19)** which is adapted to receive at least part of a data-set containing data **(e.g., column 27, lines 34-35)**, to delineate the data into multiple partitions **(e.g., column 4, lines 50-54)**, to transfer the partitions to the initial set of storage devices in accordance with the logical addresses **(e.g., column 7, lines 50-55)**, to redistribute the logical addresses among the storage devices in the extended set so as to cause a portion of the logical addresses and the partitions stored thereat to be transferred from the storage devices in the initial set to the additional storage device **(e.g., column 28, lines 40-49)**, while maintaining the balanced access

and without requiring a substantial transfer of the logical addresses among the storage devices in the initial set." **(e.g., column 9, lines 4-12; column 23, lines 49-55).**

7. *In regard to claim 31 Bridge teaches:*

"Data distribution apparatus **(e.g., Fig. 19)**, comprising: an initial set of storage devices among which are distributed logical addresses so as to provide a balanced access to the devices;" **(e.g., see abstract; column 8, lines 38-39).**

"a depleted set of storage devices, formed by subtracting a surplus storage device from the initial set;" **(e.g., column 5, lines 6-7).**

"and a processor which is adapted to receive at least part of a data-set containing data **(e.g., column 27, line 35)**, to delineate the data into multiple partitions **(e.g., column 4, lines 50-54)**, to transfer the partitions to the initial set of storage devices in accordance with the logical addresses **(e.g., column 7, lines 50-55)**, to redistribute the logical addresses and the partitions stored thereat of the surplus storage device among the storage devices in the depleted set while maintaining the balanced access and without requiring a substantial transfer of the logical addresses among the storage devices in the depleted set." **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

8. *In regard to claim 35 Bridge teaches:*

"A method **for** storing data **(e.g., abstract)**, comprising:"

"coupling one or more mass-storage devices to store partitions of data at respective first ranges of logical addresses (LAs);" **(e.g., column 5, 56-62; lines column 9, lines 36-43; Fig. 3).**

"configuring a plurality of interim devices to operate independently of one another;" **(e.g. Disk Drives 106-114 in Fig. 1)**

"assigning each interim device a respective second range of the LAs;" **(e.g., column 9, lines 4-10; extent B in disk drive 214 in Figs. 2A-2B).**

"coupling each interim device to receive the partitions of data from and provide the partitions of data to the one or more mass-storage devices having LAs within the respective second range;" **(e.g., column 9, lines 4-10; extent B in disk drive 214 in Figs. 2A-2B).**

"receiving input/output (IO) requests from host processors;" **(e.g., column 27, lines 28-37).**

"identifying specified partitions of data in response to the IO requests;" **(e.g., column 5, lines 56-62; column 9, lines 36-38).**

"converting the IO requests to converted-IO-requests directed to specified LAs in response to the specified partitions of data;" **(e.g., column 9, lines 36-42).**

"and directing all the converted-IO-requests to the interim device to which the specified LAs are assigned." **(e.g., column 9, lines 44-50).**

9. In regard to claim 39 Bridge teaches:

"A method for data distribution **(e.g., abstract)**, comprising:"

"receiving at least part of a data-set containing data;" **(e.g., column 27, line 35).**

"delineating the data into multiple equal size partitions;" **(e.g., column 4, lines 50-54).**

"transferring the partitions to an initial set of storage devices so as to provide a balanced access to the devices;" **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

"adding an additional storage device to the initial set **(e.g., column 22, lines 29-32; Fig. 16)**, thus forming an extended set of the storage devices comprising the initial set and the additional storage device;" **(e.g., column 31, lines 1-5; claim 32).**

"and redistributing the partitions among the storage devices in the extended set so as to cause a portion of the partitions to be transferred from the storage devices in the initial set to the additional storage device, while maintaining the balanced access and without requiring a substantial transfer of the partitions among the storage devices in the initial set." **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

10. In regard to claim 40 Bridge teaches:

"A method for data distribution **(e.g., abstract)**, comprising:"

"receiving at least part of a data-set containing data;" **(e.g., column 27, line 35).**

"delineating the data into multiple equal size partitions;" **(e.g., column 3, lines 24-30).**

"transferring the partitions to an initial set of storage devices so as to provide a balanced access to the devices;" **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

"removing a surplus storage device from the initial set, thus forming a depleted set of the storage devices comprising the initial set less the surplus storage device;" **(e.g., column 5, lines 6-7).**

"and redistributing the partitions stored in the surplus device to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of the partitions among the storage devices in the depleted set." **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

11. In regard to claim 41 Bridge teaches:

"Data distribution apparatus **(e.g., Fig. 19)**, comprising: an initial set of storage devices;"

"an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device;"

"and a processor which is adapted to receive at least part of a data-set containing data, to delineate the data into multiple equal size partitions **(e.g., column 3, lines 24-30)**, to transfer the partitions to the initial set of storage devices so as to provide a balanced access to the initial set of storage devices **(e.g., column 7, lines 50-55; column 23, lines 49-55)**, to redistribute the partitions among the storage devices in the extended set so as to cause a portion the partitions stored in the initial set to be transferred to the additional storage device, while maintaining the balanced access and without requiring

a substantial transfer of the partitions among the storage devices in the initial set." **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

12. In regard to claim 42 Bridge teaches:

"Data distribution apparatus **(e.g., Fig. 19)**, comprising:"

"an initial set of storage devices;" **(e.g., column 15, lines 60-61).** For example initial disk group.

"a depleted set of storage devices, formed by subtracting a surplus storage device from the initial set;" **(e.g., column 5, lines 6-7).**

"and a processor which is adapted to receive at least part of a data-set containing data **(e.g., column 27, line 35)**, to delineate the data into multiple equal size partitions **(e.g., column 4, lines 50-54)**, to transfer the partitions to the initial set of storage devices so as to provide a balanced access to the initial set of storage devices **(e.g., column 7, lines 50-55; column 23, lines 49-55)**, to redistribute the partitions of the surplus storage device among the storage devices in the depleted set while maintaining the balanced access and without requiring a substantial transfer of the partitions among the storage devices in the depleted set." **(e.g., column 7, lines 50-55; column 23, lines 49-55).**

13. In regard to claims 2 and 21 Bridge teaches:

"wherein the multiple partitions comprise equal size partitions." **(e.g., column 3, line 30).**

14. In regard to claims 3, 9, 13, 22, 28, and 32 Bridge teaches:

"wherein the data-sets comprise data from at least one of a file, file meta-data, a storage object, a data packet, a video tape, a music track, an image, a database record, contents of a logical unit, and an email." **(e.g., see column 6, lines 39-46).**

15. In regard to claims 11 and 15 Bridge teaches:

"wherein distributing the logical addresses comprises: generating a first plurality of sets of logical addresses, and wherein delineating the data comprises: assigning the at least part of the data-set a number chosen from a first plurality of different numbers;" **(e.g., column 9, lines 36-50; 3).**

"and assigning each partition among the multiple partitions a sequential partition number, and wherein transferring the partitions comprises **(e.g., column 31, lines 4-6)**: storing each partition at one of the sets of logical addresses in accordance with the sequential partition number of the partition and the number." **(e.g., column 9, lines 4-16; Fig. 2B).** *For example address of logical volume include extents A-D which is further partitioned to sequential partitions A0:A3 to D0:D3.*

16. In regard to claims 18 and 37 Bridge teaches:

"wherein the one or more mass-storage devices are coupled to provide a balanced access to the first ranges of LAs." **(e.g., column 5, lines 56-62; column 9, lines 36-43; Fig. 3).** For example directories in each "root disk group" contain information including logical volume addresses and pointer that identifies the extents or partitions.

17. In regard to claims 30 and 34 Bridge teaches:

"wherein the logical addresses comprise a plurality of sets of logical addresses, and wherein the processor **(Processor 1904 in Fig. 19)** is adapted to: assign the at least part of the data-set a number chosen from a plurality of different numbers, assign each partition among the multiple partitions a sequential partition number, and store each partition at one of the sets of logical addresses in accordance with the sequential partition number of the partition and the number." **(e.g., column 9, lines 4-16; Fig. 2B).** *For example address of logical volume include extents A-D which is further partitioned to sequential partitions A0:A3 to D0:D3.*

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> *(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.*

Claims 5, 10, 14, 19, 24, 29, 33, and 38 are rejected under 35 U.S.C. 103(a) as being unpatentable over Bridge in view of U.S. Patent Publication No. 2002/0099797 A1 to Merrell et al. (hereinafter Merrell).

18. In regard to claims 5, 10, 14, 19, 24, 29, 33, and 38 Bridge teaches all limitations included in the base claims but does not teach: "wherein the one or more storage

devices are operative in at least one of a storage area network, a network attached storage system, and an object storage architecture."

Merrell teaches: "wherein the one or more storage devices are operative in at least one of a storage area network, a network attached storage system, and an object storage architecture." **(e.g., paragraph 10 in pages 1 to 2)** *for using Storage Area Network (SAN).*

Disclosures by Bridge and Merrell are analogous because both related to storage Systems.

At the time of invention it would have been obvious to a person of ordinary skill in art to modify the storage system taught by Bridge to be utilized as the Storage Area Network disclosed by Merrell.

The motivation for using the SAN as taught by paragraph 10, pages 1-2 of Merrell is to enhance security and efficiency of the file storage system.

Therefore, it would have been obvious to combine teaching of Merrell with Bridge to obtain the invention as specified in the claims.

Claims 6-7 and 25-26 are rejected under 35 U.S.C. 103(a) as being unpatentable over Bridge in view of U.S. Patent Publication No. 2003/0005256 A1 to Grossman et al. (hereinafter Grossman).

19. In regard to claims 6 and 25 Bridge teaches all limitations included in the base claims but does not expressly teach: "wherein the number is chosen by a randomizing function."

Grossman teaches: "wherein the number is chosen by a randomizing function." **(e.g., see claim 10 in page 5)** *for generating Short-Quasi-Unique-Identifier (SQUID) randomly for identifying object in memory.*

Disclosures by Bridge and Grossman are analogous because both related to the pointers or means to identify data or objects in storage or memory devices.
At the time of invention it would have been obvious to a person of ordinary skill in art to modify the storage system taught by Bridge to include the random generation of SQUID (identifier) disclosed by Grossman.
The motivation for using the SQUID as taught by paragraph 18, page 2 of Grossman is to provide a memory object identifier that are shorter than UID, comprising only a small number of bits, while still providing similar functionality.

Therefore, it would have been obvious to combine teaching of Grossman with Bridge to obtain the invention as specified in the claim.

20. In regard to claims 7 and 26 Bridge teaches all limitations included in the base claims but does not expressly teach: "wherein the number is chosen by a consistent hashing function."

Grossman teaches: "wherein the number is chosen by a consistent hashing function." **(e.g., see claim 11 in page 5)** *for generating Short-Quasi-Unique-Identifier*

(SQUID) by hashing for identifying object in memory. The motivation for combining is based on the same rational given in rejection of claims 6 and 25.

Claims 17 and 36 are rejected under 35 U.S.C. 103(a) as being unpatentable over Bridge in view of U.S. Patent Publication No. 2003/0221063 A1 to Eguchi et al. (hereinafter Eguchi).

21. *In regard to claims 17 and 36 Bridge teaches all limitations included in the base claims but does not expressly teach:* "wherein at least one of the mass-storage devices has a slow access time, and wherein at least one of the interim devices has a fast access time."

Eguchi teaches: "wherein at least one of the mass-storage devices has a slow access time **(e.g., HDD ID 1 in Fig. 8)**, and wherein at least one of the interim devices has a fast access time." **(e.g., see paragraph 62 in page 5; HDD K in Fig. 8)** *for determining the performance of the parity groups based on the performance of physical drive set.*

Disclosures by Bridge and Eguchi are analogous because both related to the storage systems.

At the time of invention it would have been obvious to a person of ordinary skill in art to modify the storage system taught by Bridge to determine the performance of group parity based physical drive set as disclosed by Eguchi.

The motivation for using the parity group based on the drive performance as taught by paragraph 17, page 2 of Eguchi is to be able to relocate data between different storage subsystems in a manner transparent to the host system.

Therefore, it would have been obvious to combine teaching of Eguchi with Bridge to obtain the invention as specified in the claim.

ALLOWABLE SUBJECT MATTER

Claims 4 and 23 are objected to as being dependent upon rejected based claims, but would be allowable if rewritten in correct and independent form including all of the limitations of the base claim and any intervening claims.

1. The primary reason for allowance of claim 4 in instant application is the combination with the inclusion of following limitations: **wherein the first plurality of groups comprises s groups each having a different integral group number between 1 and s, wherein the number comprises an integer r randomly chosen from and including integers between 0 and s-1, wherein the sequential partition number comprises a positive integer p, and wherein the group number of the assigned specific group is (r+p)modulo(s) if (r+p)modulo(s).noteq.0, and s if (r+p)modulo(s)=0.**

: IMPORTANT NOTE :

I f the applicant should choose to rewrite the independent claims to include the limitations recited in either one of the claims, the applicant is encouraged to ***amend the***

title of the invention *such that it is descriptive of the invention as claimed as required be sec.* ***606.01*** *of the* ***MPEP****. Furthermore, the* ***summary of invention*** *and the* ***abstract*** *should be amended to bring them into harmony with the allowed claims as required by paragraph 2 of* ***sec. 1302.01*** *of the* ***MPEP****.*

As allowable subject matter has been indicated, applicant's response must either comply with all formal requirements or specifically traverse each requirement not compiled with. See ***37 C.F.R. § 1.111(b)*** *and* ***§ 707.07(a) of the M.P.E.P.***

Response to Remarks

In view of the Applicant's amendment, the objection to the abstract is withdrawn with this Office Action. The Applicant's argument with respect to the rejection of claims has carefully been considered but is not persuasive.

In regard to rejection of the independent claim 1, Applicant repeats the limitations included in the independent claim 1 (page 6 of the remarks) and then states the following:

"Claim 16 relates to a data storage system that includes: one or more mass-storage devices, coupled to store partitions of data at respective first ranges of logical addresses (LAs); a plurality of interim devices, configured to operate independently of one another, each interim device being assigned a respective second ravage of the LAs and coupled to receive the partitions of data from and provide the partitions of data to the one or more mass-storage devices having LAs within the respective second range; and one or more interfaces, which ,are adapted to receive input/output (IO) requests from host processors, to identify specified partitions of data in response to the IO requests, to convert the IO requests to converted-IO-requests directed to specified LAs in response to the specified partitions of data, and to direct all the converted-IO requests to the interim device to which the specified LAs are assigned. The Examiner asserts that Bridge discloses the claimed interface. However, although the word interface appears here in the specification of Bridge, the meaning is completely different. Bridge says: "A logical volume is the basis of the storage interface presented

to a client application 130 of storage system 100". In stark contrast, in the present invention, interfaces in the present invention comprise the stripe-cache mapping 228, and is illustrated in figure 14 as interfaces 226. An interface according to the present receive input/output (IO) requests from host processors, to identify specified partitions of data in response to the IO requests, to convert the IO requests to converted-IO-requests directed to specified LAs in response to the specified partitions of data, and to direct all the converted-IO-requests to the interim device to which the specified LAs are assigned. The Examiner asserts that this feature is disclosed at column 5, lines 56-62, and column 9, lines 36-36, of Bridge. However, the cited section merely relates to a "root disk group" that can be used to store information describing the entire disk system. The root disk group information identifies all directories that are used to maintain the disk system. In addition, it contains information about logical volumes, disk groups, and physical disk drives in disk system 102 Disk headers in the root disk group contain addresses of logical volume directories. In contrast, the present invention is completely different precisely because the interface distributes the data only into the caches, and does not control what happens later on in the disks. Then the caches distribute into the disks that are under their responsibility. In the present invention the caches are configured to operate independently of one another. There is no central database that controls what happens in all disks, but rather each cache controls what happens in it's jurisdiction. Bridge discloses a "root-disk-group" that can be used to store information describing the entire disk system, which is avoided in the present invention. Bridge's system lacks a truly distributed architecture and lacks a mapping similar to that of element 228. In sum, the entire apparatus is different, because in Bridge the distribution is done directly over the disks, and in a centralized manner, whereas the present invention provides a distributed system with a two-part distribution scheme that enables devolution of centralized control. Therefore, since Bridge does not identically disclose or suggest all of the features of claims 1 and 16, Bridge does not anticipate claims 1 and 16." (Page 6 to page 8 of Applicant's Remarks, emphasis added).

In regard to rejection of claim 16 and one or more interfaces, the Applicant argues that interface in Bridge is different from the interface in the instant application. To describe the difference Applicant states:

Applicant makes references to the specification to describe how his invention is different from disclosure by Bridge. In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which

applicant relies (i.e., interfaces in the present invention comprise the stripe-cache mapping ...) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See In re Van Geuns, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

However, the Examiner believes that Bridge teaches all limitations as recited in claim 16, Bridge further teaches:

"A logical volume manager (also called a logical disk manager) can be used to manage storage systems containing multiple disk drives. The logical volume manager configures a pool of disk drives into logical volumes (also called logical disks) so that applications and users interface with logical volumes instead of directly accessing physical disk drives. One advantage of using a logical volume manager is that a logical volume may span multiple physical disks, but is accessed transparently as if it were a single disk drive. These logical volumes appear to other components of the computer system as ordinary physical disk drives, but with performance and reliability characteristics that are different from underlying disk drives. The logical volume manager divides a physical disk drive into one or more partitions (also known as extents or subdisks). Each logical volume is composed of one or more partitions and each partition is typically defined by an offset and length." *(Column 1, lines 20-39 of Bridge; emphasis added).*

"Each logical volume is described by a list of its pieces. Each entry in the list gives the disk drive and the piece within the disk drive for one piece of the logical volume. The list itself may be stored in one or more pieces from one or more disk drives. The list forms an array ordered by the logical volume address so that it is easy to find the piece on the physical disk drive corresponding to a logical volume piece. These lists are used to translate logical volume requests into physical disk drive I/O requests. Each disk drive contains an allocation table with one entry for each of its pieces. The allocation table describes the parameters of each of the pieces on the disk drive." *(Column 5, lines 26-37 of Bridge; emphasis added).*

"In an embodiment, each logical volume is defined by a list or directory that identifies each logical volume's parameters and extents. FIG. 3 shows one embodiment of such a directory, which is termed a "logical volume directory" 302. The directory itself can be a logical volume that is stored in one or more extents from one or more disk drives. This directory is used to translate logical I/O requests into disk drive I/O requests. Referring to FIG. 3, logical volume directory 302 contains one entry for every logical volume in the disk system. Each entry contains information about the logical volume

such as, for example, its size, striping factor, stripe unit, allocation, disk group, extent size, and redundancy algorithm. The logical volume directory entry also contains one or more extent pointers." *(Column 9, lines 35-50 of Bridge; emphasis added).*

As has be shown the Bridge teaches one or more I/O interfaces that receive the logical volume I/O requests and translate (e.g., convert) the logical volume (e.g., LAs) to disk physical volume using one or more directories. In regard to rejections of claim 1 and the rejections of the other independent claims, the Applicant does not make any specific argument. He only states that the other independent claims have similar limitations and the same reasoning applies. In regard to the rejection of dependent claims has stated that since they depend from allowable independent claims therefore they are also allowable. However, as has been noted above Bridge teaches the limitations recited in the claims and therefore the Examiner maintains his position.

Conclusion

THIS ACTION IS MADE FINAL. *Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).*

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication should be directed to Hashem Farrokh whose telephone number is (571) 272-4193. The examiner can normally be reached Monday-Friday from ***8:00 AM to 5:00 PM.***

If attempt to reach the above noted Examiner by telephone are unsuccessful, the examiner's supervisor, Mr. Donald A Sparks, can be reached on (571) 272-4201. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published application may be obtained from either private PAIR or Public PAIR. Status information for unpublished application is available through Private PAIR only. For more information about PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBS) at 866-217-9197 (toll-free).

HF

HF

2007-04-12

Brian R. Peugh
Primary Examiner